arXiv:1403.2295v1 [cs.LG] 10 Mar 2014

Sublinear Models for Graphs

Brijnesh J. Jain
Technische Universität Berlin
Berlin, Germany
e-mail: brijnesh.jain@gmail.com

This contribution extends linear models for feature vectors to sublinear models for graphs and analyzes their properties. The results are (i) a geometric interpretation of sublinear classifiers, (ii) a generic learning rule based on the principle of empirical risk minimization, (iii) a convergence theorem for the margin perceptron in the sublinearly separable case, and (iv) the VC-dimension of sublinear functions. Empirical results on graph data show that sublinear models on graphs have similar properties as linear models for feature vectors.

1 Introduction

Linear classifiers and the k-nearest neighbor rule are two simple but powerful and well-investigated methods. A linear classifier is computationally efficient, makes strong assumptions about structure and yields stable but possibly inaccurate predictions [6]. In addition, linear classifiers, for example in the guise of the perceptron, form an elementary building block for more powerful methods such as deep learning [9] and other neural network architectures [8]. In contrast, the k-nearest neighbor rule is computationally less efficient, makes mild assumptions about structure and its predictions are often accurate but can be unstable [6].

While the k-nearest neighbor rule is applicable to any distance space, linear classifiers are confined to inner product spaces. The impact of linear classifiers in pattern recognition gives rise to the following questions: (1) How can we extend linear classifiers to domains other than vector spaces? (2) Suppose we can answer the first question. How can we efficiently learn such a classifier and what are its basic properties?

There are two main research directions related to the wider context of both questions. The first approach embeds patterns into a vector space and then applies linear classifiers [14, 15, 12]. The other approach defines a suitable positive semi-definite kernel for support vector learning [4, 7]. None of these directions provides an answer to either of both questions.

In this contribution, we extend linear models from feature vectors to graphs. Given a suitable graph similarity function, we extend linear functions to functions of the form $f(X) = \mathrm{sim}(W, X) + b$, where X is an input graph, W is a weight graph, and $b \in \mathbb{R}$ is the bias. We show that functions $f(X)$ on graphs correspond to piecewise linear function in some Euclidean vector space and provide a geometrical interpretation of the decision surface separating the class regions. Since those piecewise linear functions are sublinear, we refer to $f(X)$ as a sublinear function on graphs. We derive a generic update rule for learning the weight graph W and present a perceptron convergence theorem for a separable sample of graphs. Using the geometrical insight, we derive the VC dimension of sublinear functions and discuss the issue of structural risk minimization for model selection. In an empirical study, we show that the margin perceptron algorithm for graphs has similar properties as known from its vectorial counterpart.

2 Sublinear Functions on Graphs

This section introduces sublinear functions and the graph domain they are defined on.

2.1 The Space of Attributed Graphs

Let $\mathbb{A} = \mathbb{R}^d$ be a set of node and edges attributes. We consider graphs of the form $X = (\mathcal{V}, \mathcal{E}, \mathcal{A})$, where $\mathcal{V}$ represents a set of vertices, $\mathcal{E}$ a set of edges, and $\mathcal{A} \subseteq \mathbb{A}$ a subset of attributes of the nodes and edges. Node attributes take the form $\boldsymbol{x}_{ii} \in \mathcal{A}$ for each node $i \in \mathcal{V}$ and edges attributes are given by $\boldsymbol{x}_{ij} \in \mathcal{A}$ for each edge $(i, j) \in \mathcal{E}$. By $\mathcal{X}_\mathcal{G}$ we denote the space of all graphs with attributes from $\mathbb{A}$.

2.2 Sublinear Dot Product

We equip the space $\mathcal{X}_\mathcal{G}$ with a similarity function derived from the dot product of vectors. Given two graphs X and Y, the desired similarity function aims at finding one-to-one correspondences between nodes of graph X and nodes of graph Y such that the sum of corresponding node and induced edge similarities is maximized.

We encode node correspondences by a binary match-matrix $\boldsymbol{M} = (m_{ij}) \in \{0, 1\}^{m \times n}$, where m and n are the number of nodes of X and Y, respectively. If node i of X is assigned to node r of Y, then $m_{ir} = 1$, and $m_{ir} = 0$ otherwise. In addition, we demand that

$$\sum_i m_{ir} \leq 1 \quad \forall r \tag{1}$$

$$\sum_r m_{ir} \leq 1 \quad \forall i \tag{2}$$

$$\sum_{ir} m_{ir} = \min(n, m). \tag{3}$$

The constraints ask for one-to-one correspondences that can not be extended to larger subset of nodes. The first constraint enforces that any node of Y corresponds to at

most one node from X. The second constraint enforces that any node of X corresponds to at most one node of Y. Finally, the last constraint requires that each node of the smaller graph corresponds to exactly one node of the larger graph. By $\mathcal{M}^{n\times m}$ we denote the set of all binary $n \times m$-match-matrices satisfying (1)-(3).

The second concept for defining our desired similarity function are node and edge similarities. For this, we use the dot product defined on the attribute set $\mathbb{A}$. Then each match-matrix $M \in \mathcal{M}^{n\times m}$ gives rise to a kernel on X and Y

$$k_M(X,Y) = \sum_{i,j,r,s} m_{ir} m_{js} \boldsymbol{x}_{ij}^T \boldsymbol{y}_{rs},$$

where $\boldsymbol{x_{ij}}^T \boldsymbol{y}_{rs}$ denotes the dot product between attribute vectors $\boldsymbol{x}_{ij}$ of X and $\boldsymbol{y}_{rs}$ of Y. By maximizing the kernel k_M over all match matrices, we obtain the sublinear dot product

$$X \cdot Y = \max_{\boldsymbol{M}\in\mathcal{M}^{n\times m}} k_M(X,Y). \tag{4}$$

The sublinear dot product extends the dot product from vectors to graphs. It is straightforward to verify that the function $f_Y(X) = X \cdot Y$ as a pointwise maximizer of dot products is sublinear, that is convex and positively homogeneous. For this reason, we call $X \cdot Y$ sublinear. Though the sublinear dot product is not linear, it shares similar geometrical properties and generalizes the concept of maximum common subgraph [10]. It can be reduced to a special case of the graph-edit distance and is widely used in different guises as a common choice of proximity measure for graphs [2, 3, 5, 17, 18].

2.3 Sublinear Functions

A sublinear function on graphs is of the form

$$f(X) = W \cdot X + b, \tag{5}$$

where W is the weight graph and $b \in \mathbb{R}$ is the bias. As usual, we assign graph X to the positive class if $f(X) \geq 0$, and to the negative class otherwise. Then equation $f(X) = 0$ defines a decision surface

$$\mathcal{H}_f = \{X \in \mathcal{X}_\mathcal{G} \,:\, f(X) = 0\} \subseteq \mathcal{X}_\mathcal{G}$$

that separates the graph space $\mathcal{X}_\mathcal{G}$ into two class regions.

3 Graph Orbifolds

For studying sublinear functions on graphs, we need a more convenient representation of the graph space. We suggest to represent attributed graphs as points of some graph orbifold as proposed by [10, 11].

3.1 Graphs as Points in Orbifolds

For the sake of mathematical convenience, we make the following assumptions: (1) attributes $\boldsymbol{x}_{ij} \in \mathcal{A}$ of edges $(i,j) \in \mathcal{E}$ are non-zero; and (2) all graphs are of order n, where n is chosen to be sufficiently large.

Assumption (1) allows us to restrict to complete graphs, where non-edges are represented by edges with attribute zero.[1] Assumption (2) can be satisfied by adding isolated vertices with attribute zero to graphs of order less than n to obtain a graph of order n. Both assumptions are purely technical without computational and limiting impact in practice [10]. Both assumptions are purely technical without computational and limiting impact in practice. In a practical setting, we neither extend graphs to a larger order nor care about the maximum order n otherwise. See [10] for more details.

Under these assumptions, a graph X is completely specified by a matrix representation $\boldsymbol{X} = (\boldsymbol{x}_{ij})$, where the elements $\boldsymbol{x}_{ij}$ are the node and edges attributes. By $\mathcal{X} = \mathbb{A}^{n \times n}$, we denote the Euclidean space of $(n \times n)$-matrices with elements from $\mathbb{A}$.

The particular form of a matrix representation of a graph depends on how we order its nodes. Permuting the order of the nodes of a graph may result in a different matrix representation. By permuting the nodes in all possible ways, we obtain the equivalence class $[\boldsymbol{X}]$ of all matrix representations of X.

Mathematically, we can describe the equivalence class $[\boldsymbol{X}]$ of matrix representations as the orbit of $\boldsymbol{X}$ under the action of a permutation group $\mathcal{G}$ on $\mathcal{X}$. Let $\mathcal{S}_n$ be the symmetric group consisting of all permutations of n elements. Each permutation $\phi \in \mathcal{S}_n$ can be regarded as a bijective mapping $\phi : \mathcal{V} \to \mathcal{V}$ that reorders the nodes. The permutation ϕ induces a mapping on the set $\mathcal{X}$ of matrix representations:

$$g_\phi : \mathcal{X} \to \mathcal{X}, \quad \boldsymbol{X} \mapsto g_\phi(\boldsymbol{X}),$$

where $g_\phi(\boldsymbol{X})$ is the matrix representation obtained from $\boldsymbol{X}$ by reordering the vertices via the permutation ϕ. Then $\mathcal{G} = \{g_\phi \,:\, \phi \in \mathcal{S}_n\}$ is a finite group acting on $\mathcal{X}$. For a matrix representation $\boldsymbol{X} \in \mathcal{X}$, the orbit of $\boldsymbol{X}$ is the set defined by

$$[\boldsymbol{X}] = \{g(\boldsymbol{X}) \,:\, g \in \mathcal{G}\}\,.$$

In the following, we use vector- rather than matrix-based notations by stacking the columns of a matrix $\boldsymbol{X}$ to its vectorized replica $\boldsymbol{x}$. Then $\boldsymbol{x}$ is a vector representation of graph X, if π projects $\boldsymbol{x}$ to X. We identify orbits $[\boldsymbol{x}]$ with graphs X and occasionally write $\boldsymbol{x} \in X$, if vector $\boldsymbol{x}$ projects to (represents) graph X.

3.2 Sublinear Dot Product

In a graph orbifold $\mathcal{X}_\mathcal{G}$, the structural dot product defined in (4) is of the form

$$\begin{aligned} X \cdot Y &= \max\left\{\boldsymbol{x}^T\boldsymbol{y} \,:\, \boldsymbol{x} \in X, \boldsymbol{y} \in Y\right\} \\ &= \max\left\{\boldsymbol{x}^T\boldsymbol{y} \,:\, \boldsymbol{x} \in X\right\} \\ &= \max\left\{\boldsymbol{x}^T\boldsymbol{y} \,:\, \boldsymbol{y} \in Y\right\}, \end{aligned}$$

[1] If the zero-vector is a valid edge-attribute, an extra dimension serving as an edge-flag can be introduced to ensure that edges have non-zero attribute.

where $\boldsymbol{y} \in Y$ in the second and $\boldsymbol{x} \in X$ in the third line are arbitrarily chosen representations. Equality follows by symmetry and the properties of the group action [10]. Of particular interest are representations $\boldsymbol{x} \in X$ and $\boldsymbol{y} \in Y$ that attain the maximum, that is $X \cdot Y = \boldsymbol{x}^T \boldsymbol{y}$. In this case, we say $\boldsymbol{x}$ and $\boldsymbol{y}$ are optimally aligned.

4 Geometrical Interpretation

Instead of studying geometrical properties of functions and sets in the graph space $\mathcal{X}_\mathcal{G}$, it is more convenient to study lifted versions of those objects in the ambient Euclidean space $\mathcal{X}$. To this end, we first define the concept of lifting. Then we apply this concept to sublinear functions $f(X)$ and decision surfaces $\mathcal{H}_f$. Finally, we present a geometrical interpretation.

4.1 Lifting Functions and Sets

Suppose that $f : \mathcal{X}_\mathcal{G} \to \mathbb{R}$ is a real-valued function defined on graphs. Then the lift f^ℓ of f is the unique function defined on the ambient Euclidean vector space $\mathcal{X}$ such that the following diagram commutes

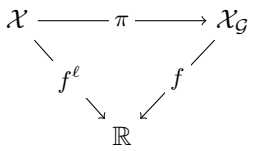

Thus, we have have $f^\ell(\boldsymbol{x}) = f(\pi(\boldsymbol{x}))$ showing that the lift f^ℓ is constant on an orbit $[\boldsymbol{x}]$, that is on all vectors $\boldsymbol{x}$ that represent the same graph $X = \pi(\boldsymbol{x})$. In a similar manner, we can define the lift of a subset $\mathcal{U}_\mathcal{G} \subseteq \mathcal{X}_\mathcal{G}$ by

$$\mathcal{U}_\mathcal{G}^\ell = \{\boldsymbol{x} \in \mathcal{X} \,:\, \pi(\boldsymbol{x}) \in \mathcal{U}_\mathcal{G}\}\,.$$

Thus, the lift of a set $\mathcal{U}_\mathcal{G}$ of graphs is the set of all vectors that represent a graph from $\mathcal{U}_\mathcal{G}$.

4.2 Lifted sublinear Functions

The lift of $f(X) = W \cdot X + b$ is a unique function defined by

$$f^\ell(\boldsymbol{x}) = \max_{\boldsymbol{w} \in W} \boldsymbol{w}^T \boldsymbol{x} + b, \tag{6}$$

where $\boldsymbol{x} \in X$. Equation (6) shows that the function $f(X)$ on graphs can be viewed as a pointwise maximizer of linear functions $h(\boldsymbol{x}) = \boldsymbol{w}^T \boldsymbol{x} + b$ in the ambient Euclidean space $\mathcal{X}$, parametrized by the different representations $\boldsymbol{w}$ of weight graph W. From eqn. (6) follows that f^ℓ is a piecewise linear function. Thus, sublinear functions on graphs correspond to piecewise linear functions in the ambient Euclidean space.

The lift of the decision surface $\mathcal{H}_f$ defined by the equation $f(X) = 0$ is of the form

$$\mathcal{H}_f^\ell = \{\boldsymbol{x} \in \mathcal{X} \,:\, f(\pi(\boldsymbol{x})) = 0\}\,.$$

Since $f(\pi(\boldsymbol{x})) = f^\ell(\boldsymbol{x})$, we see that the lifted surface $\mathcal{H}_f^\ell$ coincides with the surface $\mathcal{H}_{f^\ell}$ defined by the lifted equation $f^\ell(\boldsymbol{x}) = 0$.

4.3 Truncation

By the properties of the structural dot product, we can express a sublinear function as

$$f(X) = \max_{\boldsymbol{x} \in X} \boldsymbol{w}^T \boldsymbol{x} + b,$$

where $\boldsymbol{w}$ is an arbitrarily chosen representation of weight graph W. To evaluate f at X, we select a representation $\boldsymbol{x}$ of X that is closest to $\boldsymbol{w}$ and then compute the linear function $\boldsymbol{w}^T \boldsymbol{x} + b$.

Suppose that $\phi : \mathcal{X}_\mathcal{G} \to \mathcal{X}$ is a map satisfying the following two properties for all graphs $X \in \mathcal{X}_\mathcal{G}$:

1. $\pi \circ \phi = \mathrm{id}$
2. $W \cdot X = \boldsymbol{w}^T \phi(X)$

The first property states that ϕ maps each graph to one of its representations. Mappings with this property are called cross sections. The second property states that the cross section ϕ selects for each graph a representation that is closest to $\boldsymbol{w} \in W$. We call ϕ a fundamental cross section along $\boldsymbol{w}$.

Note that ϕ is not uniquely determined, because there are graphs that may have several different representations satisfying both properties of a fundamental cross section. As an immediate consequence, the image $\phi(\mathcal{X}_\mathcal{G})$ is also not unique, however, the closure of $\phi(\mathcal{X}_\mathcal{G})$ is well-defined and of the form

$$\mathcal{D}_{\boldsymbol{w}} = \left\{\boldsymbol{x} \in \mathcal{X} \,:\, \boldsymbol{w}^T \boldsymbol{x} \geq \widetilde{\boldsymbol{w}}^T \boldsymbol{x},\, \widetilde{\boldsymbol{w}} \in W\right\}.$$

We call the closure $\mathcal{D}_{\boldsymbol{w}}$ the Dirichlet (fundamental) domain of $\boldsymbol{w}$. A Dirichlet domain is a convex polyhedral cone [11] with the following properties:

1. $\boldsymbol{x} \in \mathcal{D}_{\boldsymbol{w}}$ iff $(\boldsymbol{x}, \boldsymbol{w})$ is an optimal alignment
2. $\pi(\mathcal{D}_{\boldsymbol{w}}) = \mathcal{X}_G$
3. π is injective on the interior of $\mathcal{D}_{\boldsymbol{w}}$

From these properties we can draw the following conclusions: (i) Since the natural projection π restricted to $\mathcal{D}_{\boldsymbol{w}}$ is surjective, each graph X has at least one representation $\boldsymbol{x}$ in $\mathcal{D}_{\boldsymbol{w}}$. (ii) Since the natural projection π is injective on the interior of $\mathcal{D}_{\boldsymbol{w}}$, different representations $\boldsymbol{x}$ and $\boldsymbol{x}'$ of the same graph lie on the border of $\mathcal{D}_{\boldsymbol{w}}$. This implies that graphs with multiple representations in $\mathcal{D}_{\boldsymbol{w}}$ form a set of Lebesgue measure zero.

From the foregoing discussion follows that studying sublinear functions f can be reduced to the study of truncated lifts f^t defined by the restriction of the lift f^ℓ to some Dirichlet domain $\mathcal{D}_{\boldsymbol{w}}$. The truncated lift f^t is a linear function

$$f^t(\boldsymbol{x}) = \boldsymbol{w}^T\boldsymbol{x} + b$$

locally defined on $\mathcal{D}_{\boldsymbol{w}}$.

4.4 Geometrical Interpretation

Linear functions $h(\boldsymbol{x}) = \boldsymbol{w}^T\boldsymbol{x} + b$ on feature vectors give rise to the following geometrical interpretation:

1. the decision surface $\mathcal{H}_{\boldsymbol{w},b}$ defined by $h(\boldsymbol{x}) = 0$ is a hyperplane;
2. the weight vector $\boldsymbol{w}$ is normal to the hyperplane $\mathcal{H}_{\boldsymbol{w},b}$;
3. $b/\left\|\boldsymbol{w}\right\|$ is the distance of the hyperplane $\mathcal{H}_{\boldsymbol{w},b}$ from the origin; and
4. $h(\boldsymbol{x})/\left\|\boldsymbol{w}\right\|$ is an algebraic measure of the distance of $\boldsymbol{x}$ from the hyperplane $\mathcal{H}_{\boldsymbol{w},b}$.

We show that these four geometrical properties carry over to sublinear functions $f(X) = W \cdot X + b$ to a certain extent.

Property 1: The decision surface $\mathcal{H}_f$ of a sublinear function f in the graph space $\mathcal{X}_\mathcal{G}$ corresponds to a piecewise linear decision surface in the ambient Euclidean space $\mathcal{X}_\mathcal{G}$. The lifted decision surface is of the form

$$\mathcal{H}_f^\ell = \left\{\boldsymbol{x} \in \mathcal{X} \,:\, f^\ell(\boldsymbol{x}) = \max_{\boldsymbol{w}\in W} \boldsymbol{w}^T\boldsymbol{x} + b = 0\right\}.$$

Since the lift f^ℓ is piecewise linear, the decision surface $\mathcal{H}_f^\ell$ is piecewise linear. We can express $\mathcal{H}_f^\ell$ as the union of hyperplane segments

$$\mathcal{H}_f^\ell = \bigcup_{\boldsymbol{w}\in W} \mathcal{H}_{\boldsymbol{w}},$$

where each hyperplane segment $\mathcal{H}_{\boldsymbol{w}}$ is defined by

$$\mathcal{H}_{\boldsymbol{w}} = \left\{\boldsymbol{x} \in \mathcal{X} \,:\, f^\ell(\boldsymbol{x}) = \boldsymbol{w}^T\boldsymbol{x} + b = 0\right\}.$$

The decision surface of the truncated lift f^t on $\mathcal{D}_{\boldsymbol{w}}$ is the intersection

$$\mathcal{D}_{\boldsymbol{w}} \cap \mathcal{H}_f^\ell$$

and coincides with the hyperplane segment $\mathcal{H}_{\boldsymbol{w}}$

Property 2: Each representation $\boldsymbol{w}$ of the weight graph W is normal to the hyperplane segment $\mathcal{H}_{\boldsymbol{w}}$.

Property 3: The distance of the decision surface $\mathcal{H}_f$ from the zero graph is given by

$$d(\mathcal{H}_f, 0) = \frac{b}{\sqrt{W \cdot W}}.$$

Observe that the distance of each linear hyperplane segment $\mathcal{H}_{\boldsymbol{w}}$ from the origin is constant with value

$$\frac{b}{\|\boldsymbol{w}\|} = \frac{b}{\sqrt{W \cdot W}},$$

because the norm $\|\boldsymbol{w}\|$ is invariant under permutations.

Property 4: The value

$$\frac{f(X)}{\sqrt{W \cdot W}} \leq d(X, \mathcal{H}_f)$$

is a lower bound of the distance $d(X, \mathcal{H}_f)$ of X from the hyperplane $\mathcal{H}_f$. Equality holds, if there is a representation $\boldsymbol{x} \in X$ and a representation $\boldsymbol{w} \in W$ such that $\boldsymbol{x}$ has an orthogonal projection onto the hyperplane segment $\mathcal{H}_{\boldsymbol{w}}$.

5 Learning

Learning typically amounts in minimizing a differentiable risk function using local gradient information. The concept of derivative, however, is unknown for functions on graphs. We solve this problem by first lifting the risk function on graphs to the Euclidean space and then applying stochastic subgradient method for minimizing the lifted risk.

5.1 The Learning Problem

The goal of learning consists in finding a weight graph W and bias b such that the sublinear discriminant $f(X) = W \cdot X + b$ minimizes the expected risk

$$E(f) = \int_{\mathcal{X}_\mathcal{G}} L(f(X), y)\, dP(X, y), \tag{7}$$

where $L(\hat{y}, y)$ is a differentiable loss function that measures the cost of predicting class $\hat{y}$ when the actual class is y and $P(X, y)$ is the joint probability distribution on $\mathcal{X}_\mathcal{G} \times \mathcal{Y}$. Since the distribution $P(X, y)$ is usually unknown, the expected risk $E(f)$ can not be computed directly. Instead, we approximate the expected risk $E(f)$ by minimizing the empirical risk

$$E_N(f) = \frac{1}{N} \sum_{i=1}^{N} L(f(X_i), y_i)$$

on the basis of N training examples $(X_i, y_i) \in \mathcal{X}_\mathcal{G} \times \mathcal{Y}$.

5.2 The Lifted Risk

The lift of the expected risk is of the form

$$E_N^\ell(f) = \frac{1}{N}\sum_{i=1}^{N} L\left(\max_{\boldsymbol{x}_i \in X_i} \boldsymbol{w}^T \boldsymbol{x}_i + b, y_i\right) \tag{8}$$

From eqn. (8) directly follows

$$E_N(f) = E_N^\ell(f) = E_N\left(f^\ell\right).$$

Since we minimize the lifted risk $E_N^\ell(f)$ as a function of $(\boldsymbol{w}, b)$, it is more convenient to express the lift f^ℓ as

$$f^\ell(\boldsymbol{w}, b|X) = \max_{\boldsymbol{x} \in X} \boldsymbol{w}^T \boldsymbol{x} + b.$$

The lift f^ℓ is convex as a function of $(\boldsymbol{w}, b)$ and non-differentiable on a subset with probability zero. Since the loss L is convex and differentiable, the same holds for the composition $L(f^\ell(\boldsymbol{w}, b|X), y)$.

5.3 Stochastic Subgradient Method

Our goal is to minimize the lifted risk $E_N^\ell(f)$ as a function of $(\boldsymbol{w}, b)$ in an iterative fashion. Each iteration t consists of choosing a training example (X_t, y_t) at random, and updating the weights $\boldsymbol{w}$ and the bias b according to the following learning rule

$$\begin{aligned} \boldsymbol{w}_{t+1} &= \boldsymbol{w}_t + \eta_t \cdot \partial_{\boldsymbol{w}_t} \\ b_{t+1} &= b_t + \eta_t \cdot \partial_{b_t}, \end{aligned}$$

where $0 < \eta_t$ is the learning rate and $\partial_t = (\partial_{\boldsymbol{w}_t}, \partial_{b_t})$ is a suitable direction along which we move the point $(\boldsymbol{w}_t, b_t)$ in order to minimize $E_N^\ell(f)$. Given the loss at iteration t

$$E_t(\boldsymbol{w}, b) = L(\hat{y}_t, y_t) = L\left(f^\ell\left(\boldsymbol{w}_t, b_t|X_t\right),\ y_t\right),$$

we define the update direction ∂_t by

$$\begin{aligned} \partial_{\boldsymbol{w}_t} &= -\nabla L(\hat{y}_t, y_t) \cdot \boldsymbol{x}_t \\ \partial_{b_t} &= -\nabla L(\hat{y}_t, y_t), \end{aligned}$$

where $\boldsymbol{x}_t$ is a representation of X_t optimally aligned with $\boldsymbol{w}_t$, and $\nabla L(\hat{y}_t, y_t)$ is the derivative of L as a function of $\hat{y}_t$.

At differentiable points $(\boldsymbol{w}_t, b_t)$ the direction ∂_t is exactly the opposite direction of the gradient of $E_t(\boldsymbol{w}, b)$. In this case, the update rule of learning a sublinear classifier on graphs coincides with its counterpart in vector spaces. At non-differentiable points $(\boldsymbol{w}_t, b_t)$ the update direction ∂_t satisfies the inequality

$$E_t(\boldsymbol{w}, b) \geq E_t(\boldsymbol{w}_t, b_t) + \partial_t^T\left((\boldsymbol{w}, b) - (\boldsymbol{w}_t, b_t)\right)$$

for all $(\boldsymbol{w}, b)$. Any vector satisfying the above inequality is a subgradient. If E_t is convex, a subgradient exists at each point $(\boldsymbol{w}, b)$. At differentiable points the subgradient of E_t is unique and coincides with its gradients.

Example (Margin Perceptron): Suppose that $\boldsymbol{x}$ and $\boldsymbol{w}$ are optimally aligned representations. The hinge loss of the margin perceptron is of the form

$$L(\boldsymbol{w}^T\boldsymbol{x} + b, y) = \max\left\{0, \lambda - y \cdot (\boldsymbol{w}^T\boldsymbol{x} + b)\right\},$$

where λ is the margin. The partial direction $\partial_{\boldsymbol{w}}$ takes the form

$$\partial_{\boldsymbol{w}} = \begin{cases} -y \cdot \boldsymbol{x} & : \text{ if } y \cdot \left(\boldsymbol{w}^T\boldsymbol{x} + b\right) \leq \lambda \\ 0 & : \text{ otherwise} \end{cases}.$$

Setting the margin to $\lambda = 0$ yields the perceptron algorithm as a special case.

5.4 Margin Perceptron Convergence Theorem

The perceptron convergence theorem states that the perceptron algorithm with constant learning rate finds a separating hyperplane, whenever the training patterns are linearly separable. We provide a weaker convergence result for the margin perceptron algorithm in the graph domain.

Suppose that $\mathcal{S} \subseteq \mathcal{X}_{\mathcal{G}} \times \mathcal{Y}$ is a sample consisting of N graphs $X \in \mathcal{X}_{\mathcal{G}}$ with corresponding class labels $y \in \mathcal{Y}$. The sample is sublinearly separable if there is a weight graph W and bias b such that

$$\operatorname{sgn}(W \cdot X + b) = y$$

for all $(X, y) \in \mathcal{S}$. We say, $\mathcal{S}$ is sublinearly separable with margin $\xi > 0$ if

$$\min_{(X,y)\in\mathcal{S}} y\,(W \cdot X + b) \geq \xi.$$

Now we can state the margin perceptron convergence theorem.

Theorem 1 *Suppose that $\mathcal{S} \subseteq \mathcal{X}_{\mathcal{G}} \times \mathcal{Y}$ is sublinearly separable with margin $\xi > 0$. Then the margin perceptron algorithm with fixed learning rate η and margin-parameter $\lambda \leq \xi$ converges to a solution (W, b) that correctly separates the sample $\mathcal{S}$ after a finite number of update steps, provided the learning rate is chosen sufficiently small.*

Proof: Suppose that $|\mathcal{S}| = N$. Since sublinear functions are convex, the sublinear dot product is convex. In addition, the composition of convex functions is convex. Hence, for each example $(X_i, y_i) \in \mathcal{S}$ the function

$$F_i(W, b) = L(W \cdot X_i + b, y_i)$$

is real-valued and convex. Then there is a positive scalar C_i that bounds the subdifferential of F_i at X_i. Suppose that

$$C = \max_{i=1,\ldots,N} C_i.$$

Let

$$F(W,b) = \sum_{i=1}^{N} F_i(W,b)$$

the risk without averaging operation, that is $F = N \cdot E_N$, where E_N is the empirical risk. Since the sample $\mathcal{S}$ is sublinearly separable by assumption, the minimum F_* of $F(W,b)$ is zero. From [13], Prop. 2.1. follows

$$\lim_{t\to\infty} F(W_t,b_t) \leq F_* + \frac{\eta \cdot C^2}{2} = \frac{\eta \cdot C^2}{2},$$

where η is the learning rate. Choosing

$$\eta \leq \frac{\xi}{C^2}$$

gives

$$\lim_{t\to\infty} F(W_t,b_t) \leq \frac{\xi}{2}.$$

Since $\xi > 0$, this implies that there is a t_0 such that

$$F_t(W_t,b_t) < \xi$$

for all $t \geq t_0$. Here, F_t refers to the example $(X_t, y_t) \in \mathcal{S}$ presented at time t. From this follows that all training examples are classified correctly after a finite number of update steps, provided that $\lambda \leq \xi$. ■

5.5 VC-Dimension

Via lifting and truncating, we can show that sublinear functions on graphs correspond to linear functions restricted to some Dirichlet domain in the ambient Euclidean space $\mathcal{X}$. From this directly follows that the VC-dimension of sublinear functions is equal to $dim(\mathcal{X}) + 1$.

In what follows, we use the generic term *item* to refer to either nodes or edges of some graph. Since graphs may have different number of items, we can control the capacity of sublinear functions by the number of items of W. By strictly applying the update rule, the number of nodes of W after learning will be equal to the number of nodes of the largest graph in the training set. A similar statement about the number of edges is not possible. If many graphs are smaller than the largest graph and if the classification problem is simple, some substructures of the weight graph W may rarely undergo an update. In this case, the classifier is too complex and we can shrink the number of items of W. Hopefully only the relevant parts of the training graphs are aligned against W. Shrinking W projects the problem to a lower-dimensional subspace.

In other cases, the classification problem is harder such that the size of the weight graph W may not capture the distribution of the pattern graphs. Then the classifier is too simple and we may increase the number of items of W, even beyond the size

of the largest graph in the training set. Due to the matching process, the training graphs will be embedded into a hopefully discriminative substructure of W. Growing W projects the problem to a higher dimensional space.

Since we can control the VC-dimension of sublinear classifiers via the size of the weight graph, we can perform structural risk minimization for model selection. This issue is out of scope. Here, the number of nodes of W is set to the number of nodes of the largest graph of the training set and no constraints are imposed on the number of edges.

6 Experiments

The goal of our empirical study is to assess the performance of the perceptron and margin perceptron algorithm for graphs.

6.1 Data

We considered the coil, letter (low, medium, high), fingerprint, and grec dataset from the IAM graph database repository [16]. All datasets are split into a predefined training, validation, and test set. For two-class problems, we selected subsets of the *coil* data set. $\text{COIL}_{2,4}$ consists of graphs representing two types of ball-shaped fruits (index 2 & 4). $\text{COIL}_{8,15}$ consists of graphs representing two types of cans (index 71 & 93). All other datasets refer to multi-class problems. The *letter* data sets consist of graphs representing distorted copies of 15 letter drawings from the Roman alphabet that consist of straight lines. The distortion levels are low, medium, and high. The *fingerprints* data consist of graphs representing fingerprint images from four classes. The *grec* data consists of graphs representing symbols from noisy versions of 22 architectural and electronic drawings.

For further details of the datasets including a description of how graphs were generated from the images, we refer to [16].

6.2 Algorithms

We compared both sublinear classifiers against the nearest neighbor method and the support vector machine with similarity kernel [1, 16].

6.3 Experimental Protocol.

We applied the graduated assignment algorithm [5] for solving graph matching problems. For multi-class problems, we wrapped both sublinear classifiers into a one-against-all classifier.

In a first step, we selected the parameters for the perceptron algorithm (learning rate η) and the large-margin perceptron (learning rate η and margin λ). For each $\eta = 0.01$, 0.05, 0.1, 0.3, 0.5, 0.7, 0.9 we trained the perceptron algorithm on the training set and evaluated the learned model on the validation set. We conducted this experiment 10

data	#(classes)	1-NN	sk-svm	perceptron		margin-perc	
				avg	max	avg	max
$COIL_{2,4}$	2	90.0	—	$95.0^{\pm 0.0}$	95.0	$95.0^{\pm 0.0}$	95.0
$COIL_{71,93}$	2	95.0	—	$96.0^{\pm 5.2}$	100.0	$98.5^{\pm 2.4}$	100.0
Letter L	15	95.6	99.2	$94.5^{\pm 0.7}$	96.0	$95.5^{\pm 0.3}$	95.7
Letter M	15	92.2	94.7	$86.1^{\pm 1.1}$	87.5	$88.7^{\pm 0.6}$	89.5
Letter H	15	83.9	92.8	$80.7^{\pm 1.1}$	82.5	$84.1^{\pm 0.5}$	84.8
F'print	4	80.1	81.7	$76.8^{\pm 1.6}$	79.1	$79.5^{\pm 2.6}$	82.4
GREC	22	97.5	(92.2)	$96.3^{\pm 0.5}$	97.0	$97.5^{\pm 0.6}$	98.1

Table 1: Classification results. Results of perceptron and margin perceptron are averaged over 10 runs. Standard deviation and maximum classification accuracy are shown for both sublinear models. Results are unknown for entries with '−'-sign.

times and selected the learning rate η_* which resulted in the best average performance on the validation set. Next, we trained the perceptron algorithm using η_* on the training and validation set. Then we evaluated the learned model on the test set. We again conducted this experiment 10 times and recorded the average accuracy, standard deviation, and the maximum accuracy. We applied the same experimental protocol for the large-margin algorithm. We adopted the optimal learning rate η_* from the perceptron algorithm and selected the optimal margin from $\lambda = 0.01$, 0.05, 0.075, 0.1, 0.125, 0.15, 0.2.

Since the 1-NN method is deterministic, we performed a single run on each data set.

6.4 Results and Discussion

Table 1 summarizes the results. Shown are the average and maximum classification accuracies as well as the standard deviations. Numbers in parentheses in the first line show the number of classes. Results of the SVM were taken form [1]. On GREC the SVM is not comparable to the other methods, because [1] used a different version than the one originally published by [16].

Generalization performance. Both perceptron algorithms perform reasonable well compared to 1-NN and SVM. Results on the letter dataset with medium and high distortion level show that the assumption of sublinear separability is too strong and yields inaccurate predictions. On average, the margin perceptron performs slightly better than the standard perceptron algorithm. In case of $COIL_{71,93}$, both perceptron algorithms show that they can learn a separating decision surface. These findings are similar to linear models in vector spaces: they are simple and powerful methods that yield possibly inaccurate results. Room for improvements is given by more extensively exploring the hyperparameters and controlling the VC-dimension via the size of the weight graph.

Computational efficiency. The graph matching problem clearly predominates the computing time of all four classifiers. Both perceptron algorithms outperform the 1-NN and SVM approach. To classify a graph, the number of graph matching problems to be solved by a perceptron algorithm in a one-against-all setting is equal to the number of classes. In a two-class problem, only a single graph matching problem needs to be solved. In contrast, nearest neighbor classifiers need to compare a graph against all training examples. The computing time of the SVM depends on the prototype selection algorithm for embedding the graphs into a vector space and by the number of support vectors. Results in [15] and follow-up publications by the same authors indicate that the SVM approach needs to compare each graph with $30\% - 60\%$ of all training examples.

7 Conclusion

This contribution generalizes linear classifiers to sublinear classifiers for graphs. To this end, we replaced the inner product of vectors by the sublinear dot product of graphs and lifted the learning problem to the ambient Euclidean space. The convergence theorem for the separable case carries over to the graph domain under mild assumptions. We can control the VC dimension via the size of the weight graph. Experiments show that perceptrons for graphs perform similarly as linear models for vectors. They are simple and efficient methods that yield possibly inaccurate predictions. Structural risk minimization for model selection and stability are prospects for further research. In addition, the stage is set for devising and analyzing further sublinear models as well as more sophisticated learning algorithms such as deep learning and other neural network architectures.

References

[1] H. Bunke and K. Riesen. Improving vector space embedding of graphs through feature selection algorithms. *Pattern Recognition* 44(9):1928–1940, 2011.

[2] T.S. Caetano, L. Cheng, Q.V. Le, and A.J. Smola. Learning graph matching. *ICCV*, 2007.

[3] T. Cour, P. Srinivasan, and J. Shi. Balanced graph matching. *NIPS*, 2006.

[4] T. Gärtner. A survey of kernels for structured data. *ACM SIGKDD*, 5(1):49–58, 2003.

[5] S. Gold and A. Rangarajan. A graduated assignment algorithm for graph matching. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 18(4):377–388, 1996.

[6] T. Hastie, R. Tibshirani, and J. Friedman. *The elements of statistical learning.* Springer New York, 2001.

[7] D. Haussler. Convolution kernels on discrete structures. *echnical report, Department of Computer Science, University of California at Santa Cruz*, 1999.

[8] S. Haykin. *Neural networks: a comprehensive foundation*. Prentice Hall, 1994.

[9] G. Hinton, S. Osindero, and Y.-W. Teh. A fast learning algorithm for deep belief nets. *Neur. Comp.*, 18(7):1527–1554, 2006.

[10] B. Jain and K. Obermayer. Structure Spaces. *The Journal of Machine Learning Research*, 10:2667–2714, 2009.

[11] B. Jain and K. Obermayer. Learning in Riemannian Orbifolds. *arXiv preprint arXiv:1204.4294*, 2012.

[12] L. Livi, A. Rizzi, and A. Sadeghian. Optimized Dissimilarity Space Embedding for Labeled Graphs. *Information Sciences*, 2014.

[13] A. Nedic and D.P. Bertsekas. Incremental subgradient methods for nondifferentiable optimization. *SIAM Journal on Optimization*, 12(1):109–138, 2001.

[14] E. Pekalska and R. Duin. *The dissimilarity representation for pattern recognition: foundations and applications*. World Scientific 64, 2005.

[15] K. Riesen, M. Neuhaus, and H. Bunke. Graph embedding in vector spaces by means of prototype selection. *Graph-Based Representations in Pattern Recognition*, p. 383–393, 2007.

[16] K. Riesen and H. Bunke. IAM graph database repository for graph based pattern recognition and machine learning. *Structural, Syntactic, and Statistical Pattern Recognition*, p. 287–297, 2008.

[17] S. Umeyama, "An eigendecomposition approach to weighted graph matching problems", *IEEE Transactions on PAMI*, 10(5):695–703, 1988.

[18] M. Van Wyk, M. Durrani, and B. Van Wyk. A RKHS interpolator-based graph matching algorithm. *IEEE Transactions on PAMI*, 24(7):988–995, 2002.